

ABSTRACT OF THE DISCLOSURE:

A multiprocessor computer system has a plurality of processing nodes which use processor state information to determine which coherent caches in the system are required to examine a coherency transaction produced by a single originating processor's storage request. A node of the computer has dynamic coherency boundaries such that the hardware uses only a subset of the total processors in a large system for a single workload at any specific point in time and can optimize the cache coherency as the supervisor software or firmware expands and contracts the number of processors which are being used to run any single workload.

Multiple instances of a node can be connected with a second level controller to create a large multiprocessor system. The node controller uses the mode bits to determine which processors must receive any given transaction that is received by the node controller. The second level controller uses the mode bits to determine which nodes must receive any given transaction that is received by the second level controller. Logical partitions are mapped to allowable physical processors. Cache coherence regions which encompass subsets of the total number of processors and caches in the system are chosen for their physical proximity. A distinct cache coherency region can be defined for each partition using a hypervisor.